Task-Specific Visualization Design

Lloyd A. Treinish
IBM Thomas J. Watson Research Center

This case study in operational weather forecasting demonstrates the principles of task-specific visualization design: defining user needs, implementing that definition, and establishing techniques for different user goals.

Visualization has matured sufficiently that user needs rather than enabling technology can drive the design of content. Improvements in technology have typically led to generalized systems as the preferred mechanism to address a diversity of visualization strategies. Such flexibility benefits research activities and applications development. However, their inherent lack of focus makes them less suitable in environments with relatively fixed tasks or user goals. This is especially true in operational situations, where users have no need to master generalized interfaces and may view their many facilities as superfluous.

Overcoming this barrier requires an understanding of user goals and how they map to visualization tasks. For example, Domik and Gutkauf modeled user needs,[1] while Card and Mackinlay[2] created a taxonomy for visualization design. Since a visualization requires domain-specific content, Jung (among others) matched both interface and composition design to task.[3] These efforts achieved the goal of decomposition after significant iteration with users.

Another approach automates the design process. For example, Zhou and Feiner discussed an expert system-based implementation focused on design elements, which also uses a taxonomy of data characteristics.[4] But this approach proves viable only with limited visualization techniques and data. For more general scientific problems, the available techniques and the diversity of the data don't lend themselves to a tractable, expert-system solution. The user has considerable domain expertise that defines the tasks and is required to interpret results. Having the user (intelligence) in the visualization process enables more effective use of that expertise and the human capacity for pattern recognition.

An alternate approach develops a set of visualization tasks coupled with appropriate designs a priori, then refines them through modest iteration. This approach employs generalized design elements and tests them to more efficiently develop focused visualizations.

Task-based visualization

To begin, consider three steps for defining visualization tasks:

1. Defining the application in terms of user needs
2. Composing design elements and interface actions to implement that definition
3. Establishing different techniques for various user goals

Prototypes help focus step 1 and converge on results for steps 2 and 3. During that refinement, the tasks are decomposed hierarchically by recognizing that the user's tasks are not the same as the visualization tasks. For example, a given user may require one or more visualization tasks, and a specific technique may support more than one user task. The desire for specific results—such as feature or event identification or communication of the results—drives the user's goals. On the other hand, the visualization tasks consist of graphical or interface actions—such as select, interact, animate, and interrogate—used for specific composition actions, like browse, analyze, or present. To test these concepts for task decomposition in visualization design requires applying them to an interesting set of problems, visualization of meteorological data for operational weather forecasting.

Related visualization work in meteorology

Visualization in meteorology predates computing—scientists drew contour maps of weather data by hand. While researchers in atmospheric sciences have been early adopters of modern 3D visualization methods, operational weather forecasting has focused on 2D visualizations. The majority of turnkey visualization systems for meteorology arose from the perspective of

0272-1716/99/$10.00 © 1999 IEEE

"one size fits all"—one interface and style of visualization independent of task to support a single class of users. Although such systems have succeeded, their focused visualizations do not address some user goals and operational efficiencies.

For mission-critical tasks, improvements in speed and effectiveness have significant impacts—the reason why weather agencies have developed focused visualization tools. One example is the Advanced Weather Interactive Processing System (AWIPS) deployed by the National Weather Service (NWS), which provides 2D visualizations.[5] A plethora of other organizations worldwide offer similar tools, termed Class I. They provide colormapped or contoured 2D scalar fields for analysis tasks by forecasters with minimal direct (graphical) interaction at a specific "layer," either the ground or a constant atmospheric pressure. Given a flat canvas for visualization design, these tools can only show a few parameters simultaneously (for example, overlaying wind as barbs or arrows, another scalar variable as line contours).

Despite the fact that operational weather centers generate large, 3D data sets, Class I techniques still dominate with few exceptions. Chief among them is Vis-5D, developed at the University of Wisconsin. It has a fixed user interface with specific visualization tools. The implementation focuses on manipulating regularly gridded data, preferably compressed to byte precision to increase the speed of operation. This yields a highly interactive tool that maps well to many meteorological data sets,[6] in use by several operational weather centers, primarily for analysis. However, for other forecasting tasks such as model assessment and dissemination, Vis-5D doesn't have the ideal interface or content.

An independent effort by Forecast Systems Laboratory (FSL) provided operational 3D visualization for analysis of weather models.[7] To eliminate duplication of the Vis-5D capabilities, this work changed direction to build directly on Vis-5D. The FSL efforts now concentrate on providing an interface consistent with other facilities (specifically AWIPS) used primarily for analysis, based on evaluation of user preferences and tasks.[8]

Fraunhofer Institute for Computer Graphics (Institut für Graphische Datenverarbeitung, or IGD) took a different approach. They've implemented independent systems focused on specific tasks. The first, Triton, was oriented toward generating 2D visualizations for nonmeteorologists.[9] The second, TriVis, is based on a related goal—providing 2D and 3D visualizations for television broadcasts.[10] The third, Rassin, provides analysis facilities directly on the native grids of meteorological data.[11]

Compositional guidelines

To enable a set of design elements useful for a variety of tasks, I have developed visualization methods within a natural coordinate system to provide a context for 3D display and interaction. These techniques provide representations of the atmosphere consistent with the data source, registered with ancillary or reference data (for example, terrain and political boundary maps). The particular task dictates the content design and choice of coordinate system to support both conceptual and physical realizations.

Since color proves critical in design, I applied knowledge of human perception via a rule-based advisory tool sensitive to the spatial frequency of data and the visualization task.[12] It serves in designing specific elements integrated into the final composition provided to users. For example, noisy data such as wind speed are primarily mapped into luminance, while smoothly varying data such as temperature are primarily mapped into opposing saturation pairs to impart an isomorphic or continuous representation.

For moisture-related data (such as humidity and precipitation), my approach combines two colormaps, mapping dry regions to brown and ranging through yellow to green for modest values. At high levels, the data map into blue, with decreasing luminance.

When using contouring to map the data onto a set of bands, my approach applies a segmented colormap with perceived ordinality. For discrete 3D representations (such as cloud surfaces), I choose uniform but complementary colors to minimize the effects of color mixing. Direct volume rendering employs the same hue, but coupled with simultaneous mapping into luminance and opacity.

I implemented several techniques for surface wind velocity, pseudo-colored by wind speed draped over a topographic surface. Using vector arrows of fixed size eliminates misleading motion cues during animation and shows gross atmospheric movement. In contrast, streamlines with directional arrows prove superior at capturing fronts, convergence zones, vortices, and so forth. On the other hand, waving flags (rigid or furled) that point in the direction of the wind have proven effective for nonmeteorologists.

Combining these approaches provided a base of techniques to present to forecasters, allowing me to spend greater effort in development rather than in progressive refinement of the visualizations. Subsequent iterations in composition were relatively minor, such as improving specific colormaps or choosing the visualization task for analysis.

Results

Task decomposition leads to three other classes of visualization: Class II (2D and 2-1/2D analysis), Class III (3D browsing), and Class IV (3D analysis). Let's look at each in turn.

Class II

You can view Class II as a superset of Class I to include 3D enhancements. It mainly supports analysis by forecasters, particularly for data comparison. Because the visualizations' appearance may be complex, it provides direct manipulation. Users may visualize up to five parameters simultaneously. These 2D variables may combine any surface or upper air layers from the same or different sources. Applying multiple techniques (such as color and height) permits illustrating them redundantly. Users can interactively select the variables and techniques independently.

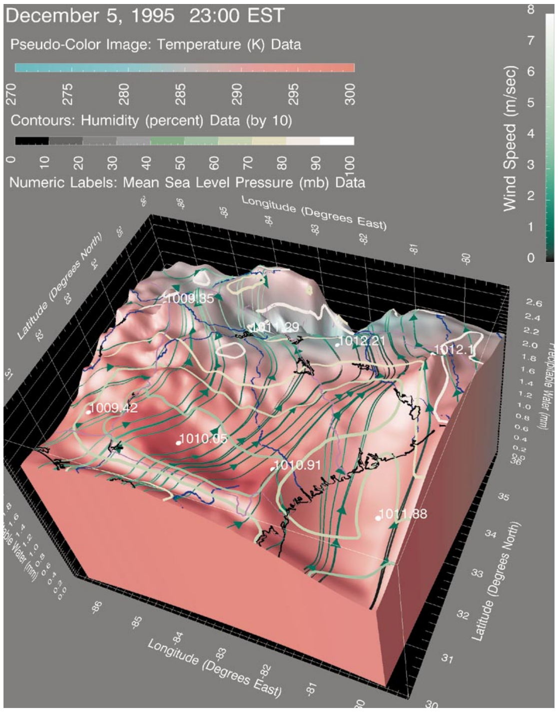

1 Class II visualization showing precipitable water as a surface.

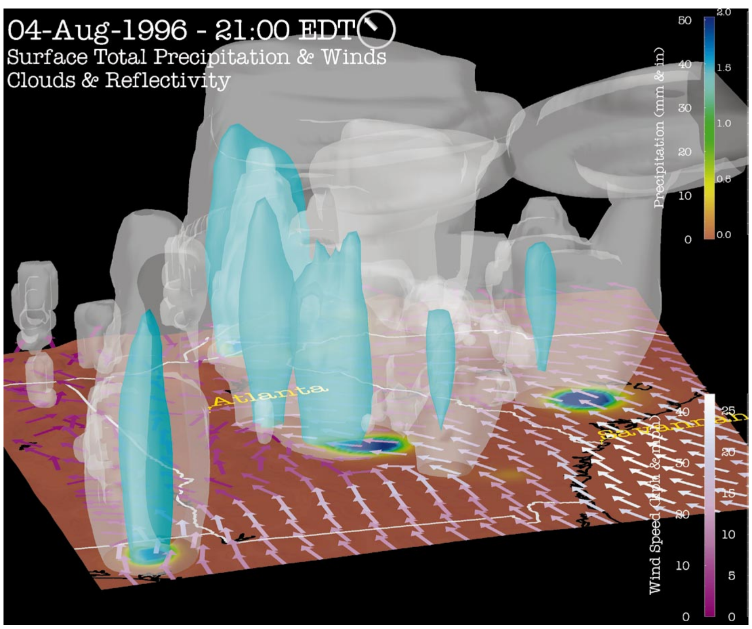

2 Class III visualization showing predicted cloud structure and rainfall in the vicinity of Atlanta during the closing ceremonies of the 1996 Olympic Games.

Figure 1 shows an example. The image represents precipitable water as the height of a shaded, deformed surface, pseudo-colored by temperature and marked with the mean sea-level pressure at discrete locations. Color-coded relative humidity contours at 10 percent intervals and streamlines of wind overlay the surface. Arrows indicate the wind direction and color, the speed. The surface is also draped with local coastline (black), state boundaries (magenta), and river (blue) maps.

Class III

Class III lets forecasters create qualitative 3D representations for both interactive investigation and animation production. The approach focuses on surface conditions and precipitation for general forecasting. The visualizations consist of a set of simplified techniques for (1) gross assessment and (2) source material suitable for public dissemination. The representations employ geographic coordinates—cartographically projected horizontally and terrain-following (true height) vertically. The techniques require high-resolution data (temporally and spatially) for coherent presentation.

Figure 2 shows an example illustrating predicted cloud structure as translucent, white isosurfaces of cloud water density at 10^{-5} kg/kg. The cloud surfaces are registered with a terrain map overlaid with coastline (black) and state (white) boundary maps, with the cities of Atlanta and Savannah marked. This representation can show atmospheric motion and potential distribution of moisture.

Pseudo-coloring the terrain by total precipitation indicates where and how much rainfall is predicted, with heavy rainfall shown as blue puddles. Translucent, cyan isosurfaces in the interior of the clouds represent forecast radar reflectivities at a threshold of 25 dBz, approximating rain shafts. The correspondence between the rain shafts and the regions of relatively heavy precipitation shows quite clearly. The topography is also overlaid with vector arrows of surface wind velocity, color-coded by speed. The visualization covers the time of the closing ceremonies of the 1996 Centennial Olympic Games in Atlanta. This visualization shows a correct prediction of thunderstorm activity in the vicinity of Atlanta, but not over the city itself.

Class IV

Class IV provides viewing and interaction tools presented in geographic coordinates—cartographically projected horizontally but at standard pressure levels vertically. This class, although similar to the visualization tasks addressed by Vis-5D and Rassin, puts greater emphasis on direct manipulation and new realization methods. Since these presentations can be visually complex even with complementary colormaps, Class IV provides facilities to interrogate and estimate data values. It also introduces the notion of a virtual metereological station, as a graphical analog for a simulated atmosphere to the type of instrumentation used to observe the real atmosphere.

Figure 3 shows an example of Class IV illustrating an analysis of atmospheric observations. A surface variable (total precipitation) appears as pseudo-color overlaid on a topographic map with rivers (blue) and coastlines (black) draped on the surface. An upper air variable (relative humidity) appears as a translucent white isosurface at 90 percent, representing a cloud boundary. Another field (temperature) is shown as a vertical slice, pseudo-color contoured.

Class IV permits visualizing any of the 3D fields available with either of these methods, surface or vertical slice. The upper-air 3D wind velocity is visualized through interactive marking of geographic locations for virtual soundings. At each location (two in this case), a vertical profile extrudes through the atmosphere. Each profile appears as a pseudo-colored tube, contoured by the variable selected for isosurface realization (here, humidity). A set of vector arrows that point in the direction of the wind shows the wind velocity along the profile. The color and length of the arrows indicate horizontal speed at these points.

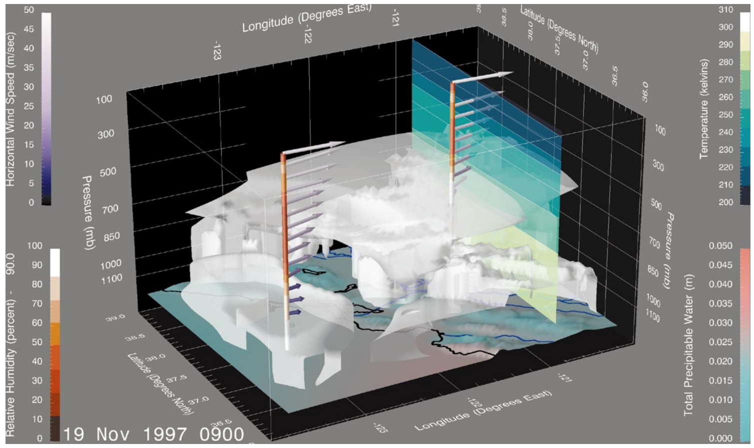

3 Class IV visualization of an analysis of atmospheric observations.

Using the visualization classes

A suite of tools for Unix and Windows workstations provides the facilities for all four visualization classes. They present simple user interfaces based on X-Window Motif for indirect interaction and OpenGL for direct 3D interaction, implemented with IBM Visualization Data Explorer (DX).[13] This suite belongs to an integrated mesoscale forecasting system called *Deep Thunder*. You can find additional information and visualization examples on the World Wide Web (http://www.research.ibm.com/weather).

The Browser application (Class III) enables model assessment. Typically, a user creates an animation with 10-minute resolution over the full model run (24 hours) after the forecaster selects the variables, techniques, and geographic view. These remain invariant throughout the animation. One or more animations generated for local playback at workstation resolution support media briefings. To aid in this selection process, the forecaster would interactively move through the geographic scene, experiment with different displays, and do limited animation either during or after model execution.

Figure 4 illustrates this process. Consider the montage of seven images, sequenced from left to right and top to bottom. Start with a 3D representation of the local area—a terrain map overlaid with state, county, coastline, and river maps, and marked with major cities. Next,

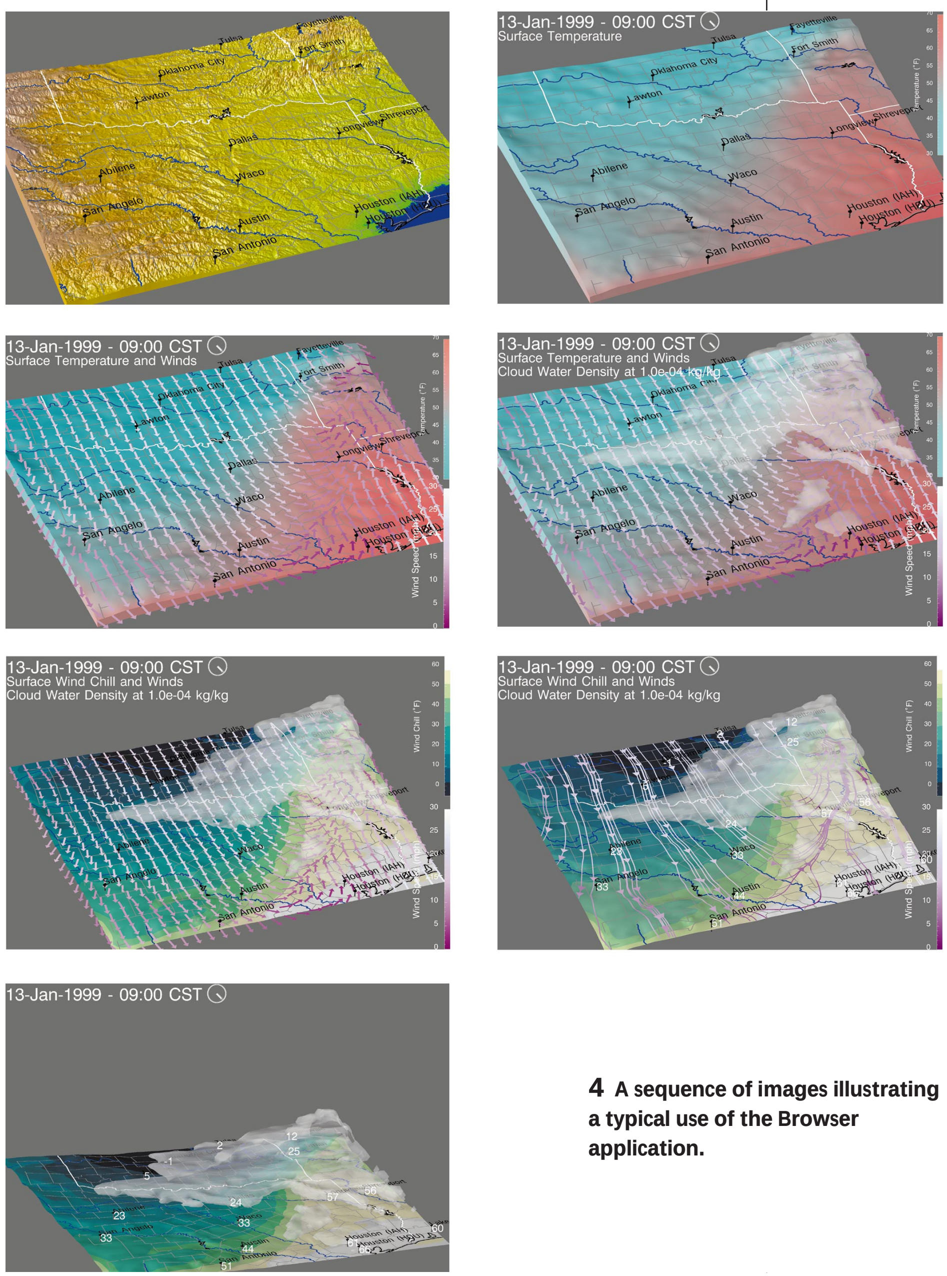

4 A sequence of images illustrating a typical use of the Browser application.

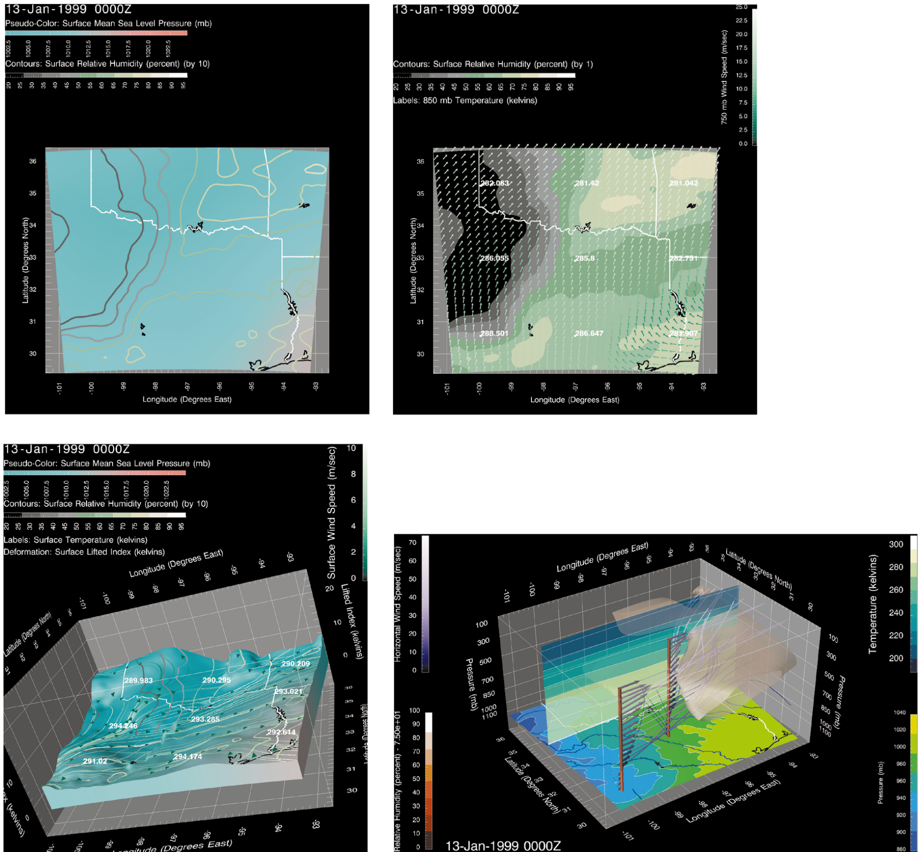

5 A sequence of images illustrating a typical use of the Slicer and Viewer applications for analysis.

overlay this topographic surface with predicted temperatures, colored by the scale at the upper right to show their continuous variation. Then, add surface winds as a set of arrows pointing in the wind direction, while the color corresponds to speed. Next, add 3D representations of predicted clouds, illustrated as a white, translucent isosurface "boundary," where the density of water exceeds a certain threshold.

From this information, a user can examine the simulation in more detail. Rather than temperature, wind chill is shown because the data indicated fairly windy conditions and low temperatures. To illustrate regions where the wind chill might be particularly low, the data appear as a set of bands, each with a distinct color increasing in value from dark to light. To examine the predicted winds in another way, they appear as colored streamlines with arrows still indicating the direction. This illustrates a front moving through the area (note the lines bunched up at the lower right).

Next, the visualization shows predicted wind chills for a specific time for major cities. From this interaction, the user could choose a representation to publish in a newspaper or on the Web to illustrate a forecast. To simplify the visualization, the wind data and some of the annotation are removed, and the geographic viewpoint of the map is changed. Now only the terrain appears, colored by bands of wind chill prediction and values at major cities along with the maps and cloud data (bottom image of Figure 4).

After each model or analysis execution, the system collects and reorganizes all of the results at hourly resolution for analysis with the Slicer and Viewer applications (Classes I, II, and IV). To understand their capabilities, consider the sequence of four images from left to right, top to bottom in Figure 5. The first image shows two surface scalar fields, mean sea-level pressure as a continuous color field, and line contours of relative humidity. Another display shows the humidity contours as color-filled bands using the same segmented colormap, but now overlaid with 850-mb temperature values at specific locations and 750-mb winds as vector arrows, colored by speed. These fields, although combined in the Class II visualization at the lower left, only show surface variables. The height of the deformed surface corresponds to lifted index, which reflects instability in the atmosphere. This representation effectively shows the motion of a front, especially in animation. The image at the lower right shows surface pressure contours combined with a 3D representation of relative humidity, temperature, and winds. The humidity data appear as an isosurface at 75 percent, corresponding to a simple cloud boundary and sampled along two vertical profiles. The temperature data are visualized as a single, vertical, contoured slice. The wind data appear as arrows using the virtual wind profiler and as streamlines using the profile points as seeds.

Conclusions and future work

Specialized interfaces and tools matched to user goals and underlying visualization tasks can provide new facilities for operational activities. I can characterize them as easy to master, even if the underlying capabilities are sophisticated. Although a generalized system can provide similar functionality, the lack of focus in its interface increases learning time. Customized systems can reduce training costs but increase the expense of development. However, a generic tool serves for both prototyping and efficient implementation by promoting high-level reuse of tools and design elements. Employing a generic toolkit (DX) also eliminated the need to implement a graphics and computational infrastructure in contrast to the low-level reuse (renderer) in efforts by FIGD or code-level modifications to a turnkey tool (Vis-5D). Since DX is built upon a unified data model that enables direct operations without transformation or compression, I did not need customization for data types and it preserved fidelity during visualization.

Unlike FSL, the work described here considers a wider variety of user goals and visualization tasks. FSL has focused primarily on interactive visualization for analysis and addressed the problem of training and usability by developing an interface consistent with existing tools.

Although FIGD has developed task-specific visualization content, they present different user interfaces and design elements, requiring additional effort in both development and training.

Class III visualizations proved more effective than expected by virtually eliminating the laborious evaluation of numerous Class I images. Conceptual models that would require inference from copious 2D data (for example, the horizontal extent of cloud dissipation in the lee of the Appalachian mountains) become obvious in 3D animations. Further, users could easily infer vertical motion based on a 3D display of clouds forming. As a result, the Browser application quickly gained favor among forecasters at the 1996 Olympic Games, enabling them to inform athletes, spectators, and officials of adverse weather conditions.

The subsequent introduction of the Slicer (Classes I and II) and Viewer (Class IV) applications into operations complements Class III. However, it uncovered problems inherent in using typical data products. Although the user could easily select data of interest, the data were not organized for interaction. In addition, not all of the variables are consistently populated, and they have incomplete metadata. This can induce user error or force increasing the application's complexity to compensate.

User-driven design has potential for other applications of precision forecasting like aviation, agriculture, broadcast, energy, and insurance. However, this requires refining the task decomposition and correlating the weather data with other information relevant to decision making.

The current applications can generate visualizations for the Web after an intermediate step of migrating the products to a Web server. This proves advantageous in an operational environment because the forecaster has content control. However, direct generation within a Web browser, which requires a simplified user interface and content, will require further refinement of the task decomposition.

The notion of task-driven customization of content and interface has succeeded in weather forecasting, but the idea also applies to other domains. Likely candidates include measurements collected from medical scanners, the output of data mining algorithms applied to relational databases, and results from terascale physics simulations. The potential benefits should encourage visualization designers to adopt these principles in their application development. ■

References

1. G.O. Domik and B. Gutkauf, "User Modeling for Adaptive Visualization Systems," *Proc. IEEE Visualization 1994*, IEEE CS Press, Los Alamitos, Calif., Oct. 1994, pp. 217-223.
2. S.K. Card and J. Mackinlay, "The Structure of the Information Design Space," *Proc. IEEE Information Visualization 1997*, IEEE CS Press, Los Alamitos, Calif., Oct. 1997, pp. 92-99.
3. V. Jung, "A System for Guiding and Training Users in the Visualization of Geographic Data," *Proc. 1st Conf. of Geo-Computation*, School of Geography, University of Leeds, Leeds, UK, Sept. 1996, pp. 470-482.
4. M.X. Zhou and S.K. Feiner, "Data Characterization for Automatically Visualizing Heterogeneous Information," *Proc. IEEE Information Visualization 1996*, IEEE CS Press, Los Alamitos, Calif., Oct. 1996, pp. 13-20.
5. National Weather Service, AWIPS Program Information, March 1998, http://www.nws.noaa.gov/msm/awips/awipsmsm.htm.
6. W.L. Hibbard et al., "Visualization of Earth and Space Science Computations," *Computer*, Vol. 27, No. 7, July 1994, pp. 65-72.
7. J.S. Snook, J.M. Cram, and J.M. Schmidt, "LAPS/RAMS," A Nonhydrostatic Mesoscale Numerical Modeling System Configured for Operational Use, *Tellus A, Dynamical Meteorology and Oceanography*, Vol. 47A, No. 5, pt. 1, Oct. 1995, pp. 864-875.
8. P.T. McCaslin, P.A. McDonald, and E.J. Szoke, "Developing and Testing a 3D VIsualization Workstation Application at FSL," *Proc. Fifteenth Int'l Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, American Meteorology Society, Boston, Jan. 1999, pp. 498-501.
9. F. Schroder, "Visualizing Meteorological Data for a Lay Audience," *IEEE Computer Graphics and Applications*, Vol. 13, No. 5, Sept. 1993, pp. 12-14.
10. F. Schroder and M. Lux, *TriVis—Professional Television Weather Presentation*, Oct. 1997, http://www.igd.fhg.de/www/igd-a4/ projects/docs/trivis/trivis_e.html.
11. M. Lux and T. Fruhauf, "A Visualization System for Operational Meteorological Use," *Proc. Sixth Int'l Conf. in Central Europe on Computer Graphics and Visualization (WSCG 98)*, University of West Bohemia Press, Plzen, Czech Republic, Feb. 1998, pp. 525-534.
12. B. Rogowitz and L. Treinish, "How Not to Lie with Visualization," *Computers in Physics*, Vol. 10, No. 3, May/June 1996, pp. 268-274.
13. G. Abram and L. Treinish, "An Extended Data-Flow Architecture for Data Analysis and Visualization," *Proc. IEEE Visualization 1995*, ACM Press, New York, Oct. 1995, pp. 263-270.

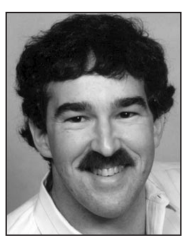

Lloyd A. Treinish *is a research staff member at the IBM Thomas J. Watson Research Center in Yorktown Heights, NY. His research interests range from visualization systems, data models, and visualization design to atmospheric and space physics. He graduated in 1978 from the Massachusetts Institute of Technology with an SM and an SB in physics, and an SB in earth and planetary sciences. He is a member of the IEEE Computer Society, ACM Siggraph, and American Geophysical Union. He is co-editor of the Visualization Viewpoints department in IEEE CG&A and was conference program co-chair for the IEEE Visualization conference in 1995, 1998, and 1999.*

Readers may contact Treinish at H4-D32, IBM Thomas J. Watson Research Center, PO Box 704, Yorktown Heights, NY 10598, e-mail lloydt@us.ibm.com.